Reducing Costs of the Modified Antarctic Mapping Mission through Automated Planning

Benjamin D. Smith Barbara E. Engelhardt Darren H. Mutz

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
{firstname.lastname}@jpl.nasa.gov

Abstract. The RadarSAT Modified Antarctic Mapping Mission (MAMM) ran from September to November 2000. The MAMM mission consisted of over 2400 synthetic aperture radar (SAR) data takes over Antarctica that had to satisfy coverage and other scientific criteria while obeying tight resource and operational constraints. Developing these plans is a time and knowledge intensive effort. It required over a work-year to manually develop a comparable plan for AMM-1, the precursor mission to MAMM. This paper describes the automated mission planning system for MAMM, which dramatically reduced mission-planning costs to just a few workweeks, and enabled rapid generation of "what-if" scenarios for evaluating mission-design trades. This latter capability informed several critical design decisions and was instrumental in accurately costing the mission.

Introduction

The Modified Antarctic Mapping Mission (MAMM) executed from September through November of 2000 onboard RadarSAT, a Canadian Space Agency (CSA) satellite. This joint NASA/CSA mission is a modified version of the First RadarSAT Antarctic Mapping Mission (AMM-1) executed in 1997. The objective of AMM-1 was to acquire complete coverage of the Antarctic continent, whereas the objective of MAMM is to acquire repeat-pass interferometry to measure ice surface velocity of the outer regions of the continent, north of 80S.

The mission objective is to perform synthetic aperture radar (SAR) mapping of the Antarctic over three consecutive 24-day repeat cycles. The SAR instrument has several "beams" each of which can be commanded to take data in rectangular swaths. These swaths are eventually compiled into a mosaic. The incidence angle of each beam is separated by a few degrees and partially overlaps the swaths of adjacent beams. The spacecraft orbit determines the swaths that are available for acquisition at any given time. The planning problem is to select a subset of the available swaths that fully cover the visible area of Antarctica and satisfy operational and resource constraints imposed by the RadarSAT Mission Management Office (MMO). The driving operational constraints are the limited on-board tape recorder (OBR) capacity and downlink opportunities, which constrain the swath subsets that will fit on the OBR between downlinks.

The AMM-1 mission demonstrated the need for an automated planning capability. The schedule for AMM-1 consisted of 850 acquisitions (swaths) over 18 days, and took over a work-year to develop manually. Despite repeated checking, this plan violated operations constraints that were not detected until the final MMO review. This inability to detect all the operations and resource constraint violations during the planning process required expensive and disruptive last-minute revisions. An automated planner could have quickly identified constraint violations, suggested repairs, and reduced the chance of errors, all of which would have significantly expedited the mission planning process.

This experience led to the development and use of an automated mission planning system for MAMM. The system expanded a set of swaths selected by the human mission planner into a detailed plan, automatically scheduled downlink activities to minimize resource costs and other criteria, and checked the resulting plan for operations constraint violations. With this system MAMM developed a 24-day mission plan containing 800 swaths in a matter of weeks, as compared to the work-year required to develop a comparable mission plan for AMM-1.

In addition to reducing the plan development effort, the MAMM planner also provided resource tracking and other details that enabled accurate costing and feasibility estimates. The MAMM planner also enabled "what-if" studies that were not possible under AMM-1. The planner quickly generated detailed variations of the baseline plan for different ground station availability assumptions. These study plans were instrumental in selecting ground stations and making other decisions about mission alternatives.

The rest of this paper describes the automated planning system that was constructed for MAMM based on the ASPEN [1, 2] planning environment. Section 2 describes the mission-planning problem, Section 3 describes the automated planning system, Section 4 describes the impact the system had on mission planning and operations.

MAMM Planning Problem

The objective of MAMM is to acquire repeat-pass SAR interferometry of Antarctica north of –80 degrees latitude to measure ice surface velocity of the outer regions of the continent. MAMM will use fine beams (high resolution) from RadarSAT to increase the accuracy of the interferometric data analysis for fast-moving glaciers found in the AMM-1 mission.

The MAMM and AMM-1 missions were conducted aboard RadarSAT. A RadarSAT acquisition plan is a time-ordered list consisting of the two parameterized activity types below. In addition to the listed parameters, each activity also has a start-time and duration parameter.

- `DataTake(beam, mode)`—Acquire SAR data in a rectangular swath along the spacecraft ground track from one of several instrument "beams" as specified by the *beam* parameter. The beams are separated by a few degrees and have overlapping ground tracks. The acquired data is either stored on the onboard recorder (OBR) or downlinked in real-time as it is acquired (RTM). The ground track varies from orbit to orbit, and repeats itself every 24 days. This means that to acquire SAR data for any given ground region the mission designer must select one opportunity from among 24 days worth of swaths on 16 different beams.
- `Downlink(station,OBR,RTM)`—Transmit data to the specified ground station. The spacecraft can transmit data from the recorder (OBR=true) and/or in real-time (RTM=true) as data is acquired. The spacecraft can transmit OBR and RTM data simultaneously to the same station. A downlink activity may only occur when the spacecraft ground track is "in view" of a ground station. There are four ground stations that can potentially receive RadarSAT data for MAMM. A downlink session can only be established with one station at a time, and only while the station is *in view* of the satellite. These requirements are specified as a list of downlink opportunities, or *masks*, each of which consists of a start time, duration, and station identifier. The opportunities vary in duration, but are always shorter than the tape capacity.

The mission-planning problem is to develop a list of datatake and downlink activities such that the SAR swaths meet the coverage and scientific requirements, the downlink activities enable all the acquired data to be transmitted, and all of the activities in the plan meet the operations constraints and resource usage constraints imposed by the RadarSAT mission management office (MMO). There are a total of 24 operational constraints, some of which are shown in Table 1. The mission planner must also try to maximize preference criteria. These include preferring ground stations that have higher reliability or lower cost, preferring RTM data to OBR (to minimize ground station and OBR resource costs), and selecting swaths to maximize coverage and other scientific criteria

Table 1: Selected Operations Constraints

OBR tapes cannot playback outside of the outer mask	Cannot transmit RTM data when recorder is in record, spin-up, or spin-down modes	OBR takes 10s to spin up, consumes 10s of tape
All recorded data must be downlinked.	Data takes shall be no less than one minute.	OBR takes 5.5s to spin down, and consumes 5.5s of tape.
OBR data can only be downlinked when a ground station outer (or inner) mask is in view	Data takes shall be at least 5.25s apart when beams are changed	OBR spin-up/spin-down between takes iff OBR data takes are > 30s apart
There will be a maximum of 6 OBR transactions per orbit	Data takes shall be at least 11s apart when beams are not changed	OBR cannot record during playback, or record during RTM data take

To determine whether the datatake and downlink activities meet the MMO constraints, the Datatake and Downlink activities must be further expanded into sub-activities. These activities consume spacecraft resources, which must also be tracked. The sub-activities are:

- Switch the SAR beam
- Start the tape recorder (takes a few seconds to spin-up)
- Stop the tape recorder (takes a few seconds to spin-down)
- Playback the entire tape for downlink, then erase it.

And the resources are:

- OBR tape used
- Number of tape transactions (on/off transitions)
- SAR instrument on-time

The swath and downlink selection decisions are tightly coupled. Minor alternations to one part of the schedule tend to require a cascade of additional changes throughout the schedule, which is part of what makes this planning problem difficult. For example, changing a swath from RTM to OBR will increase the tape usage, which may invalidate a previously selected downlink opportunity (if its duration is shorter than the recorded data volume to be downlinked). A different (longer) downlink opportunity must be chosen, and that choice may require changes to other data acquisitions, and those changes may impact other downlinks, and so on throughout the schedule.

The Planning Process

The mission planning process consists of the following four steps:

1. Select SAR swaths that cover the desired target regions in Antarctica and satisfy other scientific requirements. The swaths are selected from all the swaths that intersect the target regions during one 24 day repeat cycle as illustrated in Figure 1. Swath selection is semi-automated by SPA, a CSA tool that propagates the spacecraft orbit and determines the available swaths.
2. Create a downlink schedule. Each image must either be downlinked in real-time to a ground station that is in view during the acquisition, or stored to the data recorder and downlinked at a later opportunity. The downlink schedule must obey resource and operations constraints (e.g., recorder capacity, station visibility, time for station lock-up and exit) and conform to a priority policy (certain stations are more reliable or have lower costs than others; resource costs make real-time takes preferable to recorded ones).
3. Check the swath-and-downlink schedule for operations and resource constraint violations. This requires expanding the schedule to include the sub-activities and resources shown above, since the constraints reference these details.
4. If violations are found, return to Step 1 and modify the selected swaths to correct the problems. Modifications

include changing the swath start time, duration, and/or beam; or selecting an alternate swath on a different orbit that covers the same target area.

This process results in an acquisition plan. The MAMM planning system automates Step 2 and 3. It is up to the mission planner to select swaths (Step 1) and modify the schedule to correct violations (Step 4), since these steps require human judgement of the science impact. Swath selection (Step 1) is partially automated by a tool called SPA, developed by the Canadian Space Agency, that identifies the available swaths by propagating the spacecraft orbit but does not consider operations or downlink constraints.

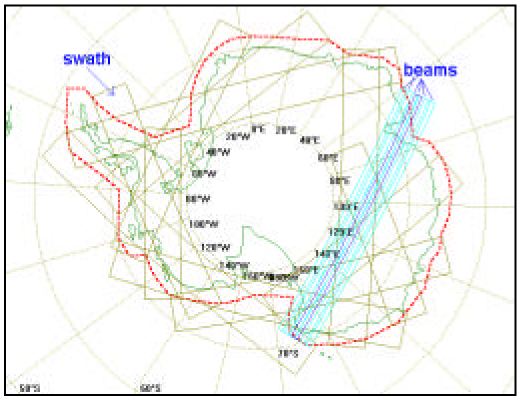

Figure 1: Swath Selection Problem. The mission planner must select swaths and beams that meet coverage and scientific criteria, and also satisfy operations constraints.

Automated Planning System

The automated planning system takes a list of SAR swaths (datatakes) selected by the mission planner, a list of accessibility masks for each ground station provided by the MMO, the station priority policy, and station mode capabilities (real-time downlink [RTM] and/or on-board recorder playback [OBR]).

The mask and swath files are combined into a single file and passed to the ASPEN planning system, which is described in more detail below. The planner expands the swaths and masks into a detailed plan that includes downlink session activities, tape on/off transitions, beam switches, and tracks resource usage. ASPEN then checks the resulting plan for operations constraints violations. The resulting plan and violations are then converted from ASPEN format to a time-ordered sequence of events and constraint violations in an Excel format that was specified by the mission planners. It also summarizes plan metrics, such as total on-board and ground station resource

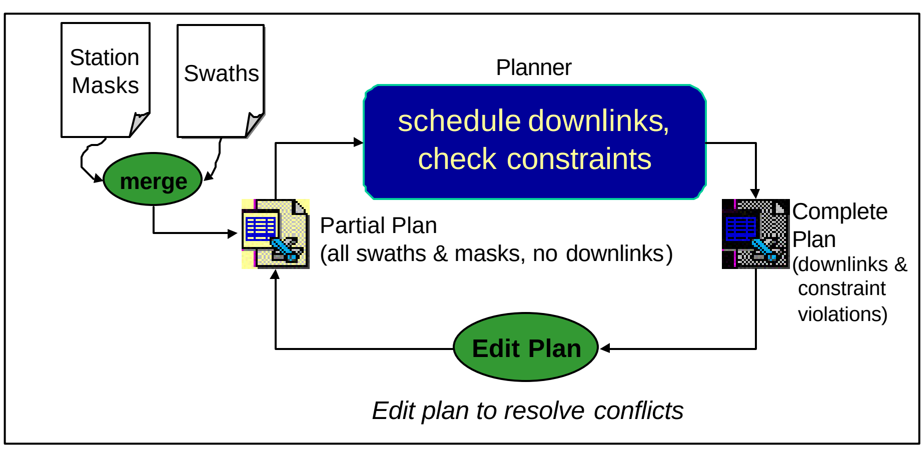

Figure 2: System Data Flow Architecture

consumption. This flow of information is documented in Figure 2.

Based on the report files, the human mission planner modifies the selected swaths as needed to resolve the conflicts or improve schedule quality. The check-and-edit cycle is repeated until a conflict-free plan is generated. This rapid feedback allows the user to generate a conflict-free plan much more quickly than is possible by hand. Maintaining the human planner in the loop enables the use of human scientific judgment in selecting swaths.

The ASPEN planner for MAMM

ASPEN [4,5] is an automated planning and scheduling system developed at the Jet Propulsion Laboratory and used for a number of space applications. Its basic operation is to find a detailed course of action—or *plan*—that achieves specified high-level goals. The goals, the actions it can take, and the operations constraints on the plan are specified in a declarative *domain model*.

The ASPEN planner has an incremental constraint tracking facility and a search facility. It uses these to process MAMM plans as follows. The search facility generates downlink activities and expands the initial plan into a detailed plan. The constraint tracker determines whether the expanded plan violates any of the constraints in the domain model. This structure is shown in Figure 3. The remainder of the section discusses this operation in more detail.

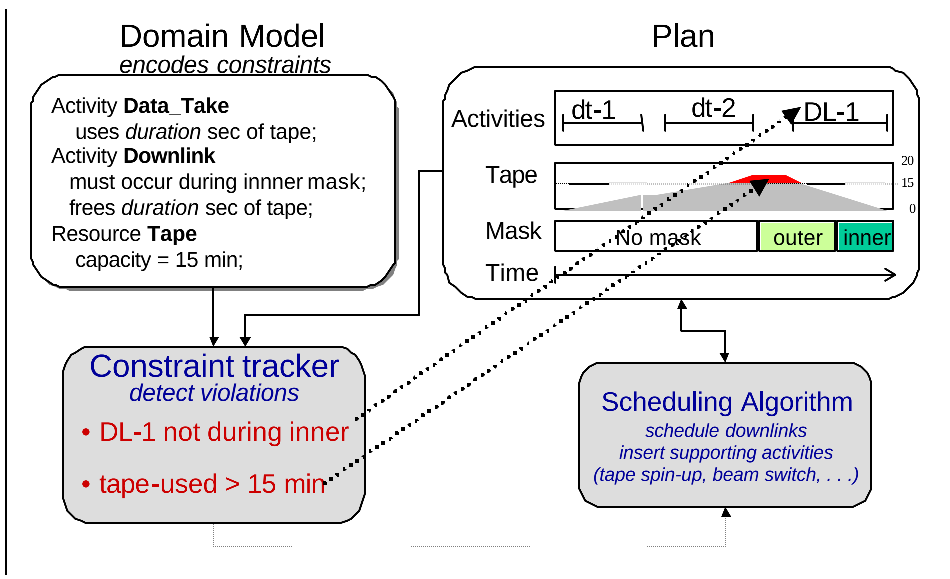

Figure 3: ASPEN planning components

The Domain Model

An ASPEN plan consists of three elements: activities, states, and resources. An activity is an action the spacecraft can perform, such as a data take or beam switch. Activities have a start time and duration and may overlap each other. A resource represents a physical or logical resource of the spacecraft, such as the onboard recorder tape or instrument on-time. A state represents a physical or logical state of the spacecraft, such as the current SAR beam or whether a given ground station is *in-view* or *not- in- view*. Each state and resource is represented as a *timeline* that shows how it evolves over time. Figure 4 shows a sample plan fragment with each of these elements.

These plan elements are related by *constraints*. These can be temporal constraints among activities (a tape spin-down must immediately follow a data take), resource constraints (a data take uses *d* seconds of OBR tape, where *d* is the duration of the data take), and state constraints (the SAR instrument must be ON during a data take). The MAMM operations constraints were encoded in terms of these constraints.

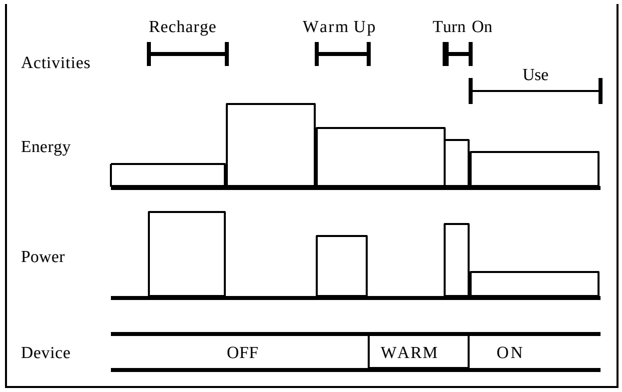

Figure 4: Timelines for activities, two resources (energy and power), and a state variable (device). Each box on the timelines is called a *timeline unit* and represents the value of that state or resource over that time period.

The Planning Algorithm

The planner takes as input a partial plan that contains just the swath activities and the downlink mask activities. The mask activities populate the mask timelines for each ground station. The planner then decides how to downlink the swaths. This downlink assignment phase assigns swaths to downlink opportunities (masks) and records these decisions in the plan by adding downlink activities for each selected mask and grounding the 'downlink mode' parameter of each swath activity to OBR or RTM accordingly.

The downlink-scheduling problem is a constrained assignment problem. Each swath must be assigned to exactly one downlink opportunity, and that assignment must satisfy OBR constraints (the duration of the selected opportunity must exceed the amount of recorded OBR data) and operations constraints (can only downlink a data take RTM if it is contained by an RTM capable mask; cannot downlink to two stations simultaneously). This problem is solved by a greedy algorithm (e.g., [3]). In each iteration it makes the best feasible assignment. If no assignment is possible, it backtracks. Since there may be no way to downlink all the selected swaths, it limits its backtracking to a two-orbit window. If no feasible solution can be found in that window, it selects a feasible schedule that downlinks the most data, and reports the lost data as a constraint violation.

After adding downlink activities and grounding the 'downlink mode' parameters of the data take activities the planner performs a limited expansion and grounding of the plan. At this point the plan consists solely of swath, mask, and downlink activities. In each iteration it selects a value for an ungrounded activity parameter, or adds an activity to satisfy an open temporal constraint. For example, if activity A is in the plan and has an open constraint that it must be before activity B, the planner will add an activity instance of type B just after activity A. At the end of this phase, the plan contains all of the activities needed to acquire and downlink the requested swaths. The resource and state timelines have also been computed based on the reservations made by the activities.

For a 24-day plan with 819 swaths and 1,068 downlink masks, the expanded plan contains 8,825 activities and over 16,000 timeline units. In a plan this size the expansion must be performed carefully to avoid unnecessary computation. The expansion uses heuristics to ensure the most computation-efficient ordering. It also uses heuristics in selecting values for grounding parameters, and for resolving disjunctive constraints. For example, when an activity is added to the schedule and imposes a resource reservation it forces all of the resource timeline units downstream of the activity to be recomputed. Placing activities in increasing time order, where possible, minimizes the computation effort.

Detecting and Reporting Constraint Violations

One the plan is expanded and grounded, the planner uses its constraint tracking mechanism to identify *conflicts*: violations of constraints in the domain model. These consist of temporal violations (e.g., data take activities are too close together), resource violations (e.g., exceeded tape capacity), and state violations.

The planner's constraint tracking facility maintains the state and resource timelines for the current plan and determines whether the constraints are satisfied. Whenever the plan changes, it incrementally recomputes the impacted timelines and constraints. State and resource timelines are computed from the state and resource constraints imposed by the activities in the current plan.

Finally, the plan and the conflicts are converted into a spreadsheet format. This is a time-ordered list of swath, mask, and downlink activities, with one row for each activity. There is one column for each resource. The value of that column for each activity (row) is the value of that resource at the end of that activity. The last column holds a list of the operations constraint violations in which that activity is involved. A table maps ASPEN conflicts to corresponding high-level operations constraints, and it is these high-level constraints that are reported in the spreadsheet.

Using Automated Planning

Before the Mission

The human mission planner used the automated planning system to generate the MAMM plan. The MAMM plan consisted of 800 swaths over 24 days. This plan was repeated on three subsequent 24-day orbit-repeat cycles to obtain interferometric SAR data pairs. The plan was not developed in one fell swoop, but rather refined over several iterations. The total development time for all five MAMM iterations was eight workweeks. Each individual plan required about two workweeks to develop. The bulk of which was spent selecting swaths and resolving conflicts. Generating downlink schedules for each draft and checking it for constraint violations was done automatically by the planning system, and required a little under an hour. By comparison, the mission plan for the first Antarctic Mapping Mission (AMM-1) consisted of 850 swaths over 18 days, and required over a work-year to develop manually.

The resource information and other details in the plan was used to evaluate draft plans, and to quickly generate and evaluate “what-if” plans for assessing alternative mission scenarios. It was particularly useful in answering the following three kinds of questions:

1. Determine the resource requirements for purposes of costing the mission and negotiating spacecraft resource allocations with the CSA.
2. How do different downlink scheduling policies impact the mission plan?
3. What is the impact of not using certain ground stations?

The ASPEN system answered each of these questions.

Question (1) was addressed using a summary page containing usage statistics generated after each analysis run of ASPEN, which gave the vital statistics for on-board recorder usage, SAR on-time, and total downlink data time broken down by station. This was invaluable in negotiating on-board resource allocations. The detailed downlink schedules were used to determine ground station cost estimates, and schedule resources. The early and detailed availability of these schedules greatly simplified this process over AMM-1.

Question (2) was addressed by performing what-if simulations using the ASPEN system. Since downlink station priorities were one of the parameters of the downlink generation algorithm, the plan was expanded and downlink schedules generated using four different possible priority systems, based on the actual cost to downlink data to certain ground stations. ASPEN supplied the data to reach a decision on the priorities and significantly impact the mission negotiations during the early stages.

Question (3) was addressed using similar what-if scenarios, where ASPEN was restricted from creating downlinks to certain stations. This enabled a closer examination of the impact of removing a ground station on the other stations and on the science collection in general. Using this information, the mission eliminated an unnecessary ground station early on in the mission operations planning phase, and saved a significant amount of funding that would have been needed to support that station during operations.

During the mission

Data takes missed due to spacecraft or ground station anomalies during the first cycle can be rescheduled for later in the cycle. The automated planner was available during operations for identifying operations conflicts in manually generated replan schedules. The system took as input the replanned schedule, and provided a list of conflicts within minutes. This capability enabled the replanning team to quickly identify and correct any constraint violations before submitting it to the MMO for a final (and more costly) check.

These capabilities were demonstrated during operations rehearsals, and were available for use in operations, but were not in fact needed during the mission. Few anomalies occurred in the first cycle, and they impacted swaths that could be trivially and confidently rescheduled manually.

Conclusions

Automated planning created a significant savings in developing mission plans, and optimized science return in a way that manual planning would take too long to perform. The planning systems also enabled rapid generation of "what-if" plans for feasibility studies, mission costing, and resource negotiations. These studies directly contributed to the quality and success of the mission, and the mission planners considered this capability an invaluable tool. Automated planning was overwhelmingly successful for MAMM, and we would expect similar successes for future missions that employ this technology.

Acknowledgements

This paper describes work performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. The authors also wish to acknowledge the support of the Canadian Space Agency in providing the RadarSAT operations constraints, and the support and sponsorship of Dr. Kim Partington during his tenure as manager of NASA's Polar Science Program.

References

1. Chien, S., Rabideau, G., Knight, R., Sherwood, R., Engelhardt, B., Nutz, D., Estlin, T., Smith, B., Fisher, F., Barrett, T., Stebbins, G., Tran, D. "ASPEN – Automating Space Mission Operations using Automated Planning and Scheduling," *SpaceOps 2000*, Toulouse, France, June2000.

2. Fukunaga, A.; Rabideau, G.; Chien, S.; and Yan, D. "Toward an Application Framework for Automated Planning and Scheduling," in *Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics, and Automation for Space*, Tokyo, Japan (1997).

3. Garey and Johnson. *Computers and Intractability: A guide to the theory of NP-completeness*. W.H. Freeman and co., New York (1979).

4. Rabideau, Knight, Chien, Fukunaga, and Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System" International Symposium on Artificial Intelligence, Robotics, and Automation for Space (ISAIRAS), Noordjwijk, The Netherlands, June 1999.

5. Sherwood et al, "ASPEN User's Guide" Internal Technical Document D-15482, Jet Propulsion Laboratory, California Institute of Technology. April, 1999.